

AMENDMENTS TO THE CLAIMS

1. (Previously presented) At least one computer readable medium encoded with instructions that, when executed by at least one processor, perform a method for generating a speech recognition model, the method comprising:
 - receiving female speech training data;
 - generating female phoneme models based on the female speech training data;
 - receiving a male speech training data;
 - generating male phoneme models based on the male speech training data;
 - determining a difference between each female phoneme model and each corresponding male phoneme model;
 - creating a gender-independent phoneme model when the difference between the compared female phoneme model and the corresponding male phoneme model is less than a predetermined value; and
 - adding, based on at least one criteria, one of the gender-independent phoneme model, or both the female phoneme model and the corresponding male phoneme model to the speech recognition model.
2. (Previously presented) The at least one computer readable medium of claim 1, wherein the at least one criteria comprises a threshold value or an upper limit for the total number of phoneme models in the speech recognition model.
3. (Previously presented) The at least one computer readable medium of claim 1, wherein determining the difference includes calculating a Kullback Leibler distance between the each female phoneme model and the each corresponding male phoneme model.
4. (Previously presented) The at least one computer readable medium of claim 3, wherein the difference is a Kullback Leibler distance quantity.

5. (Previously presented) The at least one computer readable medium of claim 1, wherein the female phoneme models, the male phoneme models, and the gender-independent phoneme models are Gaussian mixture models.

6. (Previously presented) A system for generating a speech recognition model, the system comprising:

- an input to receive speech training data; and

- a computer processor coupled to the input, the computer processor configured to:

 - receive a first set of speech training data, the first set of speech training data originating from

 - a first set of common entities;

 - generate first phoneme models based on the first set of speech training data;

 - receive a second set of speech training data, the second set of speech training data

 - originating from a second set of common entities;

 - generate second phoneme models based on the second set of speech training data;

 - determine a difference between each first phoneme model and each corresponding second phoneme model;

 - create an independent phoneme model when the difference between the compared each first phoneme model and each corresponding second phoneme model is less than a predetermined value;

and

 - add, based upon at least one criteria, one of the independent phoneme model, or both the first phoneme model and the corresponding second phoneme model to the speech recognition model.

7. (Previously presented) The system of claim 6, wherein the at least one criteria comprises a threshold value or an upper limit for the total number of phoneme models in the speech recognition model.

8. (Previously presented) The system of claim 6, wherein the computer processor is further configured to calculate a Kullback Leibler distance between the each first phoneme model and the each corresponding second phoneme model.
9. (Previously presented) The system of claim 8, wherein the difference is a Kullback Leibler distance quantity.
10. (Previously presented) The system of claim 6, wherein the first phoneme models, the second phoneme models, and the independent phoneme models are Gaussian mixture models.
11. (Previously presented) A computer program product embodied in computer memory comprising:
 - computer readable program codes executable on a computer system for generating a speech recognition model, the computer readable program codes configured to cause the system to:
 - receive a first set of speech training data, the first set of speech training data originating from a first set of common entities;
 - generate first phoneme models based on the first set of speech training data;
 - receive a second set of speech training data, the second set of speech training data originating from a second set of common entities;
 - generate second phoneme models based on the second set of speech training data;
 - determine a difference between each first phoneme model and each second phoneme model;
 - create an independent phoneme model when the difference between the each first phoneme model and the each corresponding second phoneme model is less than a predetermined value; and
 - add, based on at least one criteria, one of the independent phoneme model, or both the first phoneme model and the corresponding second phoneme model to the speech recognition model.
12. (Previously presented) The computer program product of claim 11, wherein the at least one criteria comprises a threshold value or an upper limit for the total number of phoneme models in the speech recognition model.

13. (Previously presented) The computer program product of claim 11, wherein the determining the difference includes calculating a Kullback Leibler distance between the each first phoneme model and the each corresponding second phoneme model.

14. (Previously presented) The computer program product of claim 13, wherein the difference is a threshold Kullback Leibler distance quantity.

15. (Previously presented) The computer program product of claim 11, wherein the first phoneme models, the second phoneme models, and the independent phoneme models are Gaussian mixture models.

16. (Cancelled)

17. (Previously presented) At least one computer readable medium encoded with instructions that, when executed by at least one processor, perform a method for recognizing speech from an audio stream originating from one of a plurality of data classes, each data class having class-dependent phoneme models, the method comprising:

receiving a current feature vector of the audio stream;

computing best estimates that the current feature vector belongs to each one of the plurality of data classes;

computing accumulated confidence values for each of the plurality of data classes that the current feature vector belongs to each one of the plurality of data classes, the confidence value for each data class of the plurality of data classes based on the current best estimate for the data class and on previous confidence values for the data class, the previous confidence values associated with previous feature vectors of the audio stream;

weighing the class-dependent phoneme models based on the accumulated confidence values;
and

recognizing the current feature vector based on the weighted class-dependent phoneme models.

18. (Previously presented) The at least one computer readable medium of claim 17, wherein computing best estimates includes estimating an a posteriori class probability for the current feature vector.

19. (Previously presented) The at least one computer readable medium of claim 17, wherein computing accumulated confidence values further comprises weighing the current confidence values more than the previous confidence values.

20. (Previously presented) The at least one computer readable medium of claim 17, the method further comprising determining if another feature vector is available for analysis.

21. (Previously presented) A system for recognizing speech data from an audio stream originating from one of a plurality of data classes, each data class having class-dependent phoneme models, the system comprising:

- a computer processor;

- a receiving module configured to receive a current feature vector of the audio stream;

- a first computing module configured to compute current best estimates that the current feature vector belongs to each one of the plurality of data classes;

- a second computing module configured to compute accumulated confidence values for each of the plurality of data classes that the current feature vector belongs to each one of the plurality of data classes, the confidence value for each data class of the plurality of data classes based on the current best estimate for the data class and on previous confidence values for the data class, the previous confidence values associated with previous feature vectors of the audio stream;

- a weighing module configured to weigh the class-dependent phoneme models based on the accumulated confidence values; and

a recognizing module configured to recognize the current feature vector based on the weighted class-dependent phoneme models.

22. (Original) The system of claim 21, wherein the first computing module is further configured to estimate an a posteriori class probability for the current feature vector.

23. (Previously presented) The system of claim 21, wherein the second computing module is further configured to weigh the current confidence values more than the previous confidence values.

24. (Previously presented) A computer program product embodied in computer memory comprising:

computer readable program codes executable on a computer system for recognizing speech data from an audio stream originating from one of a plurality of data classes, each data class having class-dependent phoneme models, the computer readable program codes configured to cause the system to:

receive a current feature vector of the audio stream;

compute best estimates that the current feature vector belongs to each one of the plurality of data classes;

compute accumulated confidence values for each of the plurality of data classes that the current feature vector belongs to each one of the plurality of data classes, the confidence value for each data class of the plurality of data classes based on the current best estimate for the data class and on previous confidence values for the data class, the previous confidence values associated with previous feature vectors of the audio stream;

weigh the class-dependent phoneme models based on the accumulated confidence values;

and

recognize the current feature vector based on the weighted class-dependent phoneme models.

25. (Previously presented) The computer program product of claim 24, wherein the program code configured to cause the system to compute the current best estimates includes program code configured to cause the system to determine an a posteriori class probability for the current feature vector.

26. (Previously presented) The computer program product of claim 24, wherein the program code configured to cause the system to compute the accumulated confidence values includes program code configured to cause the system to weigh the current confidence values more than the previous confidence values.

27. (Previously presented) The computer program product of claim 24, further comprising program code configured to cause the system to determine if another feature vector is available for analysis.